CHAPTER 8

Manage the Database and System Data

These topics explain how to manage the data that is used by Prime Network, including how to change how often data is archived and purged, how Prime Network organizes data in the database (embedded or external), and how to keep the system stable.

To change database passwords, see Change Password for Database Schemas, page 11-9.

For more information on the Event Archive and Fault Database and the flow of events through Prime Network, see Control Event Monitoring, page 9-1.

Overview of the Prime Network Database and Schemas

Prime Network supports two types of database installations:

Type of DB	Description	Is the DB backed up with Prime Network?
Embedded database	The Oracle database is fully integrated with Prime Network. You can use native tools to manage and monitor your database	Yes. An embedded database is backed up depending on the database profile selected during gateway installation. For more information, see Back Up and Restore Process, page 2-9.
External database	The database is managed separately from Prime Network using the tools provided by Oracle.	No. Only Prime Network registry and other data stored on the gateway is backed up. You must back up the database using the Oracle tools.

Both types of databases can be installed on the gateway server or on a separate server.

Database Schemas

A Prime Network application operating system account is created when Prime Network is installed. When Prime Network creates the database schemas, it uses this operating system account name as the default for naming the schemas.

Table 8-1 lists the database schemas that are created by Prime Network. It also provides examples of what the schema names would be if *pnuser* (the operating system account for the Prime Network application) was defined as **pn310** at installation time. You can also create the schemas manually, using different names, as described in the *Cisco Prime Network 3.10 Installation Guide*, but the purpose of each schema remains the same.

Table 8-1 Prime Network Database Schemas

Default Schema Names	Description	Example Schema Name
pnuser	Prime Network main schema that contains most Prime Network data. This schema also contains the Fault Database (the tables that are related to the fault application): • Network fault and event tables—NETWORKEVENT, ALARM, and TICKET tables. Each of these tables contain an active partition and time-based archive partitions. Tickets can be manually or automatically archived. When data is archived, it is moved to an archive partition based on the object timestamp. Archive partitions which exceeds the history size (14 days by default) are deleted. • Non-network fault and event tables—SYSTEMEVENT, AUDITEVENT, SECURITYEVENT, PROVISIONINGEVENT, NEWTRAPEVENT, NEWTRAPVALUE, MONITOR. Each of these tables contain only time-based archive partitions. Archive partitions which exceed the history size (14 days by default), are deleted. To change the settings that control when events are purged from the Fault Database, see Events, Alarms, and Tickets, page 8-8. A special process also purges tickets when the number of tickets exceeds its threshold; see Manage an Embedded Database, page 8-15.	**pn310**
*pnuser*_**ep**	Prime Network Event Archive (event persistence) schema that contains the following tables: • HP_SYSLOG • HP_TRAP The *pnuser*_**ep** schema contains all raw events that are sent from devices to Prime Network. Specifically, these raw events are received by the Event Collector (on AVM 100) and persisted in the Event Archive. The data is arranged in time-based partitions. When a partition's age exceeds its history size (14 days by default), the data is deleted. You can disable event archiving using the procedure in Disable Saving Raw Events to the Event Archive, page 8-12. To change the settings that control when events are purged from the Event Archive, see Events, Alarms, and Tickets, page 8-8.	**pn310_ep**
*pnuser*_**rep**	Prime Network reports schema that contains synonyms based on the *pnuser* schema tables; it is used by the reports mechanism. Reports are deleted according to the workflowEngine integrity test; see Reports, page 8-13.	**pn310_rep**
*pnuser*_**ep_rep**	Prime Network reports schema that contains synonyms based on the *pnuser*_**ep** schema tables; it used by the reports mechanism. Reports are deleted according to the workflowEngine integrity test; see Reports, page 8-13.	**pn310_ep_rep**

Table 8-1 Prime Network Database Schemas (continued)

Default Schema Names	Description	Example Schema Name
*pnuser*_**dwe**	Prime Network Workflow Engine schema that contains all data (templates and workflows) related to the Workflow Engine. The Workflow Engine is described in Manage Workflows, Command Scripts, and Activations, page 10-1. Workflows are deleted according to the workflowEngine integrity test; see Table 8-4 on page 8-7.	**pn310_dwe**
*pnuser*_**xmp**	Prime Network Change and Configuration Management schema that contains all data related to the functions of Change and Configuration Management. This feature is optionally installed with Prime Network and is described in the *Cisco Prime Network 3.10 Installation Guide*. For more information on Change and Configuration Management, see the *Cisco Prime Network 3.10 User Guide*.	**pn310_xmp**
*pnuser*_**admin**	User with database administrator permissions who can run maintenance tasks—such as gathering statistics—on the other Prime Network database schemas. If this user is created with the proper permissions (as described in the installation guide), Prime Network will run a cron job called **every_24_hours.cmd** that gathers statistics on other database tables. This provides an automatic method for generating database statistics, which is recommended for better performance. For more information, see the *Cisco Prime Network 3.10 Installation Guide*.	**pn310_admin**

Table 8-2 lists the main tables used by the *pnuser* schema. They are listed here to aid in troubleshooting purposes (for example, if you see any missing statistics messages in the AVM 11 log (*NETWORK_HOME*/logs/11.out), they could be referring to tables that are no longer used by Prime Network).

Table 8-2 Database Tables in the pnuser Schema

Primary Tables (Tables with More Traffic)		
ALARM	NEWTRAPEVENT	SECURITYEVENT
AUDITEVENT	NEWTRAPVALUE	SEVERITYASPECT
NETWORKEVENT	NOTIFICATION	SYSTEMEVENT
NEWAFFECTEDSNC	PROVISIONINGEVENT	TICKET

Table 8-2 Database Tables in the pnuser Schema (continued)

Primary Tables (Tables with More Traffic)		
Tables with Less Traffic		
BOSRESULTS	MAP	RECONCILIATIONASPECT
BOSUSER	MAPASPECT	REPORT
BUSINESSOBJECT	MAPDATAASPECT	REPORTDATA
CLIENTREGISTRY	MARTINITUNNELPEER	SCOPE
CONNECTIONTP	NETWORKPSEUDOWIRE	SERVICE
CUSTOMER	NETWORKVLAN	SITE
EFPCROSSCONNECT	NOTE	SWITCHINGENTITY
ETHERNETFLOWDOMAIN	OIDARRAYS	VIRTUALROUTER
ETHERNETSERVICE	PASSWORDHISTORYENTRY	VLANPERSISTEDLINK
ETHFLOWPOINT	PERMISSION	VPLSFORWARD
EVC	PSEUDOWIREEDGE	VPLSINSTANCE
HIERARCHYNODE	PWSWITCHINGENTITY	VPN
LCA		

Control How Data is Saved, Archived, and Purged

The Prime Network defaults for saving and deleting data ensure that current data remains available, while not impacting system performance. Table 8-3 lists the defaults for removing data from the system. You can adjust these settings according to the needs of your deployment. These mechanisms are described in the following topics:

- What is the Difference Between Clearing, Archiving, and Purging Fault Data?, page 8-5
- How the Data Purging Mechanism Works, page 8-6
- Events, Alarms, and Tickets, page 8-8
- Configuration Archives and Software Images, page 8-13
- Jobs, page 8-13
- Reports, page 8-13
- Diagnostic Data (Graphs Tool), page 8-15
- Workflows and Activations, page 8-14

Table 8-3 lists the default settings for deleting data from Prime Network.

Table 8-3 Default Settings for Purging Data

Data	Purged After (Default):
Database—Fault Database[1]	14 days
Database—Event Archive	14 days
Jobs	Never purged

Table 8-3 Default Settings for Purging Data

Data	Purged After (Default):
Reports	90 days
Prime Network backups for systems with external database	5 backups
Prime Network backups for systems with embedded database[2]	16 backups
Executed activations	7 days
Executed workflows	7 days
Diagnostics (Graphs) tool	29 days
Configuration Archive files and change logs	30 days
Software Images	n/a (manual deletions only)

[1] Tickets are deleted 14 days after they are moved to an archive partition. Events in an archive partition are deleted 14 after they are received by the system. For more information, see What is the Difference Between Clearing, Archiving, and Purging Fault Data?, page 8-5.

[2] See Manage an Embedded Database, page 8-15 for information on additional checks that are performed by Prime Network.

What is the Difference Between Clearing, Archiving, and Purging Fault Data?

Table 8-3 lists the default settings for when data is purged (deleted) from Prime Network, either from the database, or from gateway file systems. Fault data—that is, the data saved in the Fault Database—has a more complex purging mechanism that also involves *clearing* and *archiving*.

Note In some cases a distinction is drawn between *network events* and *non-network events*. Network events are Service, Trap, and Syslog events. Non-network events are System, Security, and Provisioning events.

Clearing

A cleared ticket means the event, alarm, or ticket should no longer be considered a problem. Users can manually clear (force clear) an item from the Vision or Events GUI clients, or by using BQL. When an item is cleared, its severity icon changes to a green check mark, providing a visual indication that the problem has been addressed. (Acknowledging an event is different. Acknowledging indicates that someone is *aware* of the issue. Acknowledging does not change the severity icon; it just changes its Acknowledged value to **True**.)

To maintain system stability, tickets that are very large (have 150 associated events) are automatically cleared so that system performance is not degraded. This is done by default.

You can configure Prime Network to automatically clear tickets based on a certain severity or age. This feature is not enabled by default. To set this auto-clear mechanism, see Change Fault Settings: Clear, Archive, and Purge Fault Data, page 8-8.

Archiving

Archived data is data that is moved to an archive partition in the Fault Database. When fault information is first saved, it is saved in the *active partition* of a database event table. Database event tables also have *archive partitions*, which are further divided into time-based subpartitions. When an event or ticket is archived, it is moved from the active partition to the archive partition. When it is moved to the archive partition, the event or ticket begins aging, based on its detection time. It is deleted after 14 days, by default.

Note Archiving network events, tickets, and alarms is *not* the same as saving data in the Event Archive. The Event Archive is the database area where *raw events* are saved; it is completely different from the archive and active partitions that contain fault data.

Only tickets, alarms, and network events (Service, Trap, and Syslog events) are archived. When a ticket is archived, all of the events and alarms associated with the ticket are also archived.

Automatic archiving is done to protect system performance. A ticket is automatically archived if it meets *any* of the following conditions:

- The ticket contains an alarm that has more than 150 events associated with it
- The ticket has been in a cleared, unchanged state for 1 hour
- The ticket exceeds the configured ticket threshold (16,000 by default)

You can also manually archive and event, alarm, or ticket from the GUI client or by using BQL.

Purging

When an event (or any other data) is purged from Prime Network, it is permanently removed from the database or other storage directories. This data can only be retrieved if it has been saved to an external location. Most data is purged after 14 days, but this depends on its type of data. For more information, see How the Data Purging Mechanism Works, page 8-6.

How the Data Purging Mechanism Works

Prime Network maintains system stability by running cron jobs to maintain the database and eliminate clutter in the system, especially fault management data. Some jobs are run every 12 hours, while others are run every hour.

Different cron jobs are run on different schedules. To check the current schedules, use this procedure.

Step 1 Using an SSH session, log into the Prime Network gateway as *pnuser*. (*pnuser* is the operating system account for the Prime Network application, created when Prime Network is installed; for example, **pn310**.)

Step 2 Use the following command to list the contents of the crontab file for user *pnuser*. The local/cron directories listed below are all located in *NETWORKHOME*.

```
# crontab -l
# Cisco Prime Network crontab file
# contains scheduled tasks for user prime-network
* * * * * if [ -f local/cron/every_1_minute.cmd ]; then local/cron/every_1_minute.cmd >
/dev/null 2>&1; fi
* * * * * /var/adm/cisco/prime-network/scripts/keep_alive_port_watchdog.pl > /dev/null
2>&1
0 * * * * if [ -f local/cron/every_1_hour.cmd ]; then local/cron/every_1_hour.cmd >
/dev/null 2>&1; fi
0 4,16 * * * if [ -f local/cron/every_12_hours.cmd ]; then local/cron/every_12_hours.cmd >
/dev/null 2>&1; fi
0 23 * * * if [ -f local/cron/every_24_hours.cmd ]; then local/cron/every_24_hours.cmd >
/dev/null 2>&1; fi
```

```
0,10,20,30,40,50 * * * *  if [ -f local/cron/every_10_minutes.cmd ]; then
local/cron/every_10_minutes.cmd > /dev/null 2>&1; fi
```

(The port watchdog script is part of the AVM protection mechanism and is described in AVM 100 and Unit Server High Availability, page 5-3.)

If desired, you can modify when the jobs run by editing the crontab file. For example, the following line in the crontab file runs the file every_12_hours.cmd at 4:00 a.m. and 4:00 p.m.:

```
0 4,16 * * * local/cron/every_12_hours.cmd > /dev/null 2>&1
```

Table 8-4 lists some of the integrity tests performed by Prime Network. These tests run on a regular basis to ensure system stability and purge old data. Prime Network performs archives and purges fault data according to the settings described in Events, Alarms, and Tickets, page 8-8.

If you have an embedded database, additional purging checks are performed as described in Manage an Embedded Database, page 8-15. These settings are defined in the registry unless otherwise noted.

Table 8-4 Integrity Tests

Test Name	Description
analyze	Generates a System event if the period between the current date and the date each database table was analyzed is larger than the analyze-Period setting. Archiving database logs is also enabled, and Prime Network saves a maximum of 10 database logs.
backup	Backs up the registry, encryption keys, and crontab files. By default, backups are saved to *NETWORKHOME*/backup. Backups are performed every 12 hours at 4:00 a.m. and 4:00 p.m. (Registry backup settings are described in Manage the Prime Network Software Image and Backups, page 2-1.)
businessObject	Checks for invalid OIDs in business objects. If more than two invalid business tags are found, Prime Network generates an event containing the list of OIDs.
capacity	Checks the disk space capacity and sends alarms. Alarms are sent when the disk capacity reaches 80% and 90%.
checkDbClock	Ensures that database clock is synchronized with the NTP server.
jobSchedulerPruning	Ensures that jobs have been deleted according to the system settings. (This setting is controlled in the Prime Network Administration GUI client; see Jobs, page 8-13).
mapAspect	Removes mapAspect OIDs which are not connected to any hierarchy.
oidArrays	Removes OIDs which exist in the OidArrays table, but not in a parent table.
reports	Deletes reports after 90 days. (This setting is controlled in the Prime Network Administration GUI client; see Reports, page 8-13).
unusableIndexes	Checks for unusable table indexes and, if found, rebuilds them.
workflowEngine	Deletes all completed or aborted workflow and activation instances after 7 days.

Events, Alarms, and Tickets

⚠
Caution Consult with your Cisco account representative before changing these settings. Making the settings smaller could result in immediate and permanent removal of fault data. Making the settings larger could result in slow data retrieval performance; the system might require additional storage and some database tuning; and backups might require more time.

This table lists the fault data management settings that you can adjust, and where to get more information.

To change the setting that controls:	See:
• When data is deleted from the Fault Database or Event Archive (default: 14 days) • When to split the database partitions in the Fault Database and the Event Archive (default: 1 hour) • Whether to auto-clear tickets, and at what thresholds (default: disabled) • When to remove events from a Vision GUI client device inventory window (to protect data retrieval performance) (default: 6 hours)	Change Fault Settings: Clear, Archive, and Purge Fault Data, page 8-8
• Auto-archiving tickets based on the *total number of ticket*s in the system (default: 16,000 tickets) • Auto-archiving tickets based on the *size* of tickets (default: Has an alarm with 150 events)	How Tickets are Auto-Archived in the Database, page 8-10
• Whether raw events are saved to the Event Archive (default: enabled)	Disable Saving Raw Events to the Event Archive, page 8-12

Change Fault Settings: Clear, Archive, and Purge Fault Data

The Event Management Settings window controls the following settings for the Fault Database and the Event Archive:

- When archived data is deleted (purged)
- When database partitions are split

Every hour Prime Network monitors the size of tables in the Fault Database, deleting old data and splitting the partitions. These settings are controlled using **Global Settings > Event Management Settings**.

All database tables contain an *active* partition and an *archive* partition. Archive partitions are further divided into time-based subpartitions. When data is archived, it is moved to an archive partition. The archive partitioning and data purging is controlled by the settings in this window.

In addition, the Inventory Event Viewer settings control when events are removed from the inventory event display Prime Network Vision. This display is launched when a user views a device's inventory in Prime Network Vision, allowing users to see some of the events and tickets on devices within their scope. By default, network and provisioning events are removed from the display after 6 hours, and no more than 15,000 events are displayed. Users can adjust this setting from their Prime Network Vision GUI client (using **Tools > Options** in their client). Changes made from the client will override the settings controlled from Administration GUI client.

Caution Consult with your Cisco account representative before changing these settings. Making the settings smaller could result in immediate and permanent removal of fault data. Making the settings larger could result in slow data retrieval performance; the system might require additional storage and some database tuning; and backups might require more time.

Step 1 Select **Global Settings > Event Management Settings** from Prime Network Administration.

Step 2 Make your desired changes to the following settings.

Field		**Description**
Fault Database	Remove events from database after ____ days	Number of days after which archived data will be deleted from each archived database partition. The default is 14.
	Database partition size (in hours)	Number of hours after which each database partition will be split. The default is 1 hour. (For database sizing guidelines and other capacity planning information, contact your Cisco account representative.)
Event Archive	Remove events from database after ____ days	Number of days after which raw event data will be deleted from each database partition. The default is 14. To disable saving any raw events to the Event Archive, see Disable Saving Raw Events to the Event Archive, page 8-12.
	Database partition size (in hours)	Number of hours after which each database partition will be split. The default is 1 hour. (For database sizing guidelines and other capacity planning information, contact your Cisco account representative.)
Inventory Event Viewer	Maximum history size (in hours)	Number of hours after which network and provisioning events are removed from the inventory event viewer in Prime Network Vision. The default is 6 hours. These settings are overridden if a user makes local changes to their Prime Network Vision GUI client (using **Tools > Options** in their client).
Auto-Clear Tickets	Automatically clear tickets	System clears the tickets that are older than a predefined time and severity. This is disabled by default.
	Severity	Select severity of the tickets (Critical, Major, Minor, Warning) that should be cleared. This is disabled by default.
	Days since last modification	Clears the ticket if the ticket was not modified for the specified number of days. This is disabled by default.

Step 3 Click **Apply**. The changes will take effect in the next partitioning process execution (which is done once an hour). You can restore the default settings at any time by clicking **Restore**.

How Tickets are Auto-Archived in the Database

Caution Consult with your Cisco account representative before changing these settings. Making the settings smaller could result in immediate and permanent removal of fault data. Making the settings larger could result in slow data retrieval performance; the system might require additional storage and some database tuning; and backups might require more time.

As described in Table 8-1 on page 8-2, the tickets table contains an active partition and time-based archive partitions. When a ticket is archived, it is moved to an archive partition based on the object timestamp. These topics describe how tickets are auto-archived.

Auto-Archiving of Cleared Tickets

Prime Network automatically archives cleared tickets which have not changed in the last hour. This setting is controlled in the registry.

Caution Consult with your Cisco account representative before changing these settings. Making the settings smaller could result in immediate and permanent removal of fault data. Making the settings larger could result in slow data retrieval performance.

Table 8-5 Registry Settings for Automatic Archiving of Cleared Tickets

Registry Entry	Description	Default Value
autoArchivingTimeout	Remove Cleared tickets that have not changed in this period of time (in milliseconds).	3600000 (1 hour)

Step 1 Log into the gateway as *pnuser* and change to the Main directory.

```
# cd $ANAHOME/Main
```

Step 2 To change the autoArchivingTimeout setting to 90 minutes:

```
# ./runRegTool.sh -gs 127.0.0.1 set 0.0.0.0 "site/plugin/AlarmPlugin/autoArchivingTimeout" 
5400000
```

Step 3 Restart AVM 11 using **networkctl**.

Auto-Archiving Based on Total Number of Tickets

Prime Network checks how many tickets are saved in the Fault Database to see if they should be archived, as follows:

- When the total number of tickets in the Fault Database exceeds 12,800, it generates a System event.
- When the total number of tickets in the Fault Database exceeds 16,000, it archives tickets in groups of 400.

Table 8-6 shows the registry parameters that control this type of Fault Database purging.

Caution Consult with your Cisco account representative before changing these settings. Making the settings smaller could result in immediate and permanent removal of fault data. Making the settings larger could result in slow data retrieval performance; the system might require additional storage and some database tuning; and backups might require more time.

For information on the format of the **runRegTool.sh** script, see Changing Registry Settings Using runRegTool.sh, page C-2.

Table 8-6 Registry Settings for Purging the Fault Database Based on Number of Tickets

Registry Entry	Description	Default Value
ticketRedThresholdAmount	When the number of open tickets surpasses this amount, archive the number of tickets specified by ticketArchivingBulk.	16000
ticketYellowThresholdPercentage	When the number of open tickets surpasses this percentage of ticketRedThresholdAmount, generate a system message.	80
wakeUpMessageInterval	Interval for checking the number of open tickets (in milliseconds).	60000 (1 minute)
ticketArchivingBulk	Alarm is generated once it crosses upper threshold after this many polling cycles.	400

If you have installed an embedded database, see the additional database management topics in Manage an Embedded Database, page 8-15.

Auto-Archiving Based on the Size of Tickets

Every five minutes, Prime Network checks the Fault Database to see if it contains any large tickets that should be archived. If large tickets accumulate, they can negatively affect system performance. A ticket is considered large if it has more than 150 events associated with an alarm. To protect system performance, Prime Network does the following:

- If a large ticket is found, it generates a System event similar to the following:

  ```
  The system contains the following XXX ticket(s) with more than 150 events per alarm.
  You can manually archive these tickets or the system will automatically archive them
  in: 15 minutes
  ```

  If the user does not respond within 15 minutes, Prime Network archives the tickets.

- If more than 1500 large tickets are found, it will send this System event:

  ```
  There are more than XXX excessively large tickets in the system (tickets with more
  than 150 events per alarm).
  ```

Table 8-7 shows the registry parameters that control this type of Fault Database purging.

Caution Consult with your Cisco account representative before changing these settings. Making the settings smaller could result in immediate and permanent removal of fault data. Making the settings larger could result in slow data retrieval performance; the system might require additional storage and some database tuning; and backups might require more time.

For information on the format of the **runRegTool.sh** script, see Changing Registry Settings Using runRegTool.sh, page C-2.

Table 8-7 Registry Settings for Purging Large Tickets From the Fault Database

Registry Entry	Description	Default Value
findLargeTicketsMessageInterval	Interval for searching for large tickets (in minutes).	5
maxTicketSize	When the number of events associated with an alarm surpasses this number, consider it a large ticket and generate a System event.	150
autoRemoveTimeInterval	Interval at which to archive a large ticket (in minutes) after sending System event.	15
oversizedTicketAmountLimit	When the number of large tickets surpasses this number, generate a System event.	1500

If you have installed an embedded database, see the additional topics in Manage an Embedded Database, page 8-15.

Disable Saving Raw Events to the Event Archive

By default, Prime Network archives all event notifications it receives from devices and saves them in the Event Archive. Events are saved according to the settings that are configured in the **Global Settings > Event Management Settings** window in the Administration GUI client (see Events, Alarms, and Tickets, page 8-8).If you do not want to save any raw events to the Event Archive, you can disable it using the following procedure.

Note If you disable this feature, the data will not be available for event-related reports.

Step 1 Log into the gateway as *pnuser* and change to the Main directory by entering the following command:

```
# cd $ANAHOME/Main
```

Step 2 Issue the following command to disable saving raw events to the Event Archive:

```
# ./runRegTool.sh -gs 127.0.0.1 set 0.0.0.0
“site/trap/agents/trap/netEventPersistencyEnabled” false
```

To reenable saving raw events to the Event Archive, use this command:

```
# ./runRegTool.sh -gs 127.0.0.1 set 0.0.0.0
"site/trap/agents/trap/netEventPersistencyEnabled" true
```

Step 3 Restart the Event Collector AVM (AVM 100) from Prime Network Administration by right-clicking the AVM and choosing **Actions > Stop** and (when it is down) **Actions > Start**.

Configuration Archives and Software Images

Prime Network Change and Configuration Management data is deleted according to these settings:

- Device configuration files and change logs are saved for 30 days by default. After that, they are deleted from the archive.
- Software image files are not deleted; they can only be manually removed using the Change and Configuration Management GUI client.

For more information, see the *Cisco Prime Network 3.10 User Guide*.

Jobs

The retention policy for job runs can be configured using the Job Manager Settings page. Old job runs which do not comply to the configured policy will be automatically purged. By default, no jobs are purged.

To set up or change Job Manager purge settings:

Step 1 Choose **Global Settings > Job Manager Settings**.

Step 2 Configure the settings that control when job runs will be purged from Prime Network.

Field	Description
Purge Job Runs After	Specifies how long to save a job run. The time is measured from when the job run is created (in days).
Store Up to	Specifies the maximum number of job runs, after which job runs should be purged. When this number is exceeded, Prime Network deletes the oldest job runs (first in, first out). Prime Network runs a purge by size check every time a new job runs is created or a user changes the settings on this page. This feature is disabled by default.

If these settings are changed to lower values, after the changes are applied, Prime Network immediately deletes all job runs that exceed the thresholds.

Step 3 Click **Apply** to immediately apply your settings.

Reports

You can run different types of reports from the Prime Network window using the Reports main menu. This feature is described in the *Cisco Prime Network 3.10 User Guide*. The Report Settings page in the Global Settings drawer controls:

- When reports should be purged. Reports are saved in the database and in a gateway file system (in an intermediate format that is rendered to HTML or PDF when viewed). By default, they are purged after 90 days. This page also shows you how much space reports are currently consuming.
- Whether users can share reports (create public reports). If a report is public, all users can view the report; public reports are *not* filtered according to scopes or security privileges.

The settings do not affect user permissions for report actions such as adding, deleting, canceling, and so forth. Users can still perform all actions on reports they create; they can view other reports only if the reports are public. Administrators are the only users who can perform all actions on all reports.

Note We recommend that you use these default settings in order to reduce system clutter. Allowing report data to accumulate could affect system performance.

To set up or change global report settings:

Step 1 Choose **Global Settings > Report Settings**.

Step 2 Configure the settings that control when reports will be purged from Prime Network, using dates, size, or both.

Field	Description
Purge report after: ___ days	Specifies how long to save a report. The time is measured from when the report is created. If you do not check this box, Prime Network defaults to 90 days. The Prime Network integrity service runs a job every 12 hours to purge all reports that exceed this age.
Store reports up to: ___ MB	Specifies the maximum disk size, in MB, at which reports should be purged. When this space setting is exceeded, Prime Network deletes the oldest reports (first in, first out). Prime Network runs a purge by size check every time a new report is created or a user changes the settings on this page. This feature is disabled by default.

If these settings are changed to lower values, after the changes are applied, Prime Network immediately deletes all reports that exceed the thresholds.

Step 3 The Enable Shared Reports check box specifies whether users can create public reports. When a report is public, all users can view the contents; reports are *not* filtered according to scopes or security privileges. Changes to this setting are applied to all subsequent new reports.

- If not selected, no users will be able to create public reports. Users will only be able to view their own reports.
- If selected, users have the option to create public reports and share them with other users.

Step 4 Click **Apply** to immediately apply your settings.

After you click **Apply**, the report settings are applied to all existing and new reports. You can restore the Prime Network default settings at any time by clicking **Restore** and **Apply**.

Workflows and Activations

All templates, workflows, and activations are saved in the database under the *pnuser*_**dwe** schema. All executed workflow and activation instances are purged according to the MAX_WORKFLOW_AGE_ENTRY. By default, this entry is set to 7, which means activation instances are deleted after 7 days. Purging is done according to the integrity test listed in Table 8-4 on page 8-7.

Templates can only deleted from the Administration GUI client or by using a BQL command. How to do this from the GUI client is described in Manage Workflow and Activation Templates, page 10-4.

Diagnostic Data (Graphs Tool)

Data gathered by the Prime Network Monitoring tool is purged after 28 days as described in Prevent System Overloads (Advanced Overload Prevention/Safe Mode), page 8-23.

Backups

Prime Network performs backups on a regular basis for Prime Network data and embedded databases. For more information, see Back Up and Restore Process, page 2-9.

Number of backups saved for Prime Network are as follows:

- Backups of embedded database: 8 days
- Backups of Prime Network data:
  - Installations with external databases: 16 backups
  - Installations with embedded databases: 16 backups

Manage an Embedded Database

Prime Network performs regular checks to ensure the health of the embedded database. Prime Network also provides native utilities for adding storage, collecting database logs and reports, and other maintenance tasks. These are all described in the following topics:

- Overview: How Prime Network Monitors an Embedded Database, page 8-15
- Embedded Database Events and Errors, page 8-16
- Stop, Start, and Change Embedded Database Settings (embdctl Utility), page 8-18
- Change the SMTP Server for Embedded Database Notifications, page 8-22

Overview: How Prime Network Monitors an Embedded Database

Prime Network performs regular maintenance checks and backups for embedded databases. Backups are enabled as part of the installation process. If you did not enable backups, you can do so using the procedure in Back Up and Restore Process, page 2-9. That topic also provides information on backup schedules, how many backups are saved, and the backup location.

Table 8-8 lists the regular maintenance checks performed by Prime Network.

Table 8-8 Cron Jobs for Maintaining the Embedded Database

Cron Job Task	Description
Monitor disk usage on database server	Hourly job that checks database disk usage (on server host) for data files, redo logs, backup files, and so on. If any directory exceeds a threshold, an e-mail and System event is sent. Event severity depends on threshold: • 50-70%—Warning event • 70-80%—Minor event • 80% and above—Major event See Database Disk Usage Alerts, page 8-16, for additional information about this problem.
Check available space in tablespaces	Hourly job that checks whether tablespaces listed in *NETWORKHOME*/Main/scripts/embedded_db/cron/TS_ALERTS.prm. If threshold is exceeded, a new data file is added to tablespace, and an e-mail and System event is sent. Event severity depends on threshold: • 80-90%—Minor event • 90% and above—Major event See Database Tablespace Usage Alerts, page 8-16, for additional information about this problem.
Check database backup log for errors	Daily job that checks backup logs for errors. Removes logs over 14 days old.
Clean database log and trace files	Hourly job that removes database log and trace files more than 31 days old.

Embedded Database Events and Errors

Prime Network monitors the embedded database and generates System events when necessary.

Database Disk Usage Alerts

Prime Network will continue to generate events (one hour later, at the next cron job) if the same directory's disk usage surpasses the *next* threshold, or a different directory's disk usage surpasses any threshold. If the disk space is unchanged, no new System events are generated.

If the problem continues:

1. Ask your system administrator to add disk space to the relevant file systems.
2. If more disk space cannot be added, contact the Cisco Technical Assistance Center for information on how to reduce history size. This will not change the disk usage, but will eliminate the need to add disk space.

Database Tablespace Usage Alerts

Note You can change the thresholds by editing the TS_ALERTS.prm file. Prime Network will use the new threshold numbers when it performs the next hourly cron job.

If a tablespace exceeds its capacity, Prime Network will add a new data file to the tablespace. Prime Network will generate an hourly system event until the problem is fixed. If the problem continues, do the following:

1. If you have the required disk space, add data files using the **add_storage_for_tablespace.pl** utility. See Add Datafiles to a Specific Tablespace (add_storage_for_tablespace.pl), page 8-21.
2. Contact the Cisco Technical Assistance Center.

Oracle Errors Monitored by Prime Network

Table 8-9 lists the Oracle errors that are monitored by Prime Network. If you receive any of the following errors, contact the Cisco Technical Assistance Center (TAC).

Table 8-9 Oracle Database Function Error Messages

Error Code and Message	Possible Reason
ORA-00600 internal error code, arguments: [*string*], [*string*], [*string*], [*string*], [*string*], [*string*], [*string*], [*string*]	This is the generic internal error number for Oracle program exceptions. This indicates that a process has encountered an exceptional condition.
ORA-00604 error occurred at recursive SQL level string	An error occurred while processing a recursive SQL statement (a statement applying to internal dictionary tables).
ORA-00050 operating system error occurred while obtaining an enqueue	Could not obtain the operating system resources necessary to cover an oracle enqueue. This is normally the result of an operating system user quota that is too low.
ORA-00052 maximum number of enqueue resources (string) *exceeded*	Ran out of enqueue resources.
ORA-00053 maximum number of enqueues exceeded	Ran out of enqueue state objects.
ORA-00055 maximum number of DML locks exceeded	Ran out of DML lock state objects.
ORA-00059 maximum number of DB_FILES exceeded	The value of the DB_FILES initialization parameter was exceeded.
ORA-00060 deadlock detected while waiting for resource	Transactions deadlocked one another while waiting for resources.
ORA-00250 archiver not started	An attempt was made to stop automatic archiving, but the archiver process was not running.
ORA-00255 error archiving log string of thread string, sequence # string	An error occurred during archiving.
ORA-00257 archiver error. Connect internal only, until freed.	The archiver process received an error while trying to archive a redo log. If the problem is not resolved soon, the database will stop executing transactions. The most likely cause of this message is the destination device is out of space to store the redo log file.
ORA-01033 ORACLE initialization or shutdown in progress	An attempt was made to log on while Oracle is being started up or shutdown.
ORA-01035 ORACLE only available to users with RESTRICTED SESSION privilege	Logins are disallowed because an instance started in restricted mode. Only users with RESTRICTED SESSION system privilege can log on.

Table 8-9 Oracle Database Function Error Messages (continued)

Error Code and Message	Possible Reason

`ORA-01110`
`data file string: (string)`
	Reports the file name. This error accompanies other errors that explain the problem associated with this file.

`ORA-01116`
`error in opening database file (string)`
	At attempt e to open a database file failed. Most likely the file is inaccessible. Accompanying errors will provide the file name.

`ORA-01520`
`number of data files to add (string) exceeds limit of string`
	CREATE TABLESPACE statement specifies more files than is permitted for this database.

`ORA-01536`
`space quota exceeded for tablespace '`*`string`*`'`
	The space quota for the segment owner in the tablespace has been exhausted and the operation attempted the creation of a new segment extent in the tablespace.

`ORA-01659`
`unable to allocate MINEXTENTS beyond string in tablespace string`
	Failed to find sufficient contiguous space to allocate MINEXTENTS for the segment being created.

`ORA-27041`
`Unable to open file`
	An attempt to open a file failed. Check the accompanying error messages for the file name.

`ORA-27100`
`shared memory realm already exists`
	Tried to start duplicate instances, or tried to restart an instance that had not been properly shutdown.

`ORA-27102`
`out of memory`
	—

`ORA-27103`
`internal error`
	—

`ORA-27146`
`post/wait initialization failed`
	OS system call failed.

Stop, Start, and Change Embedded Database Settings (embdctl Utility)

Note If you are using gateway server high availability, freeze the cluster services *before* using **emdbctl** with the **stop**, **start**, **restore**, **restore_db**, or **enable_backup** options. These options will stop and restart the cluster services. If the cluster is running and detects that the services are down, it may attempt to restart them. When used with Oracle ADG, reconfigure the database replication after restoring the primary DB. For more information on replication process, see *Cisco Prime Network 3.10 Gateway High Availability Guide*.

Use the **emdbctl** command to perform embedded database backup and restore operations, collect logs and reports, and other administrative actions. The **emdbctl** command is located in *NETWORKHOME*/Main/scripts/embedded_db. It takes the following options:

emdbctl [**--start** | **--stop** | **--enable_backup** | **--backup** | **--restore** | **--restore_db**| **--collect** | **--change_backup_time** | **--set_smtp_server**]

Option	Description	See:
--stop	Stops Prime Network on the gateway and units, and stops the embedded database services and listener.	This topic for examples.
--start	Starts the embedded database services and listener, and starts Prime Network on the gateway and units (if the units are down).	
--enable_backup	Enables the automatic backup mechanism.	Enable Embedded Database Backups, page 2-14
--backup	Backs up the embedded database and Prime Network, including the registry.	Back Up and Restore Process, page 2-9
--restore	Restores the embedded database *and* Prime Network, including the registry using valid backup files.	Back Up and Restore Process, page 2-9
--restore_db	Restores the embedded database only.	
--collect	Collects embedded database logs and reports. It collects logs and trace files from the database server, runs a database diagnostic tool, zips the output together, and copies it to the gateway at *NETWORKHOME*/Main/logs/emdb/ana_collector.zip. It can be run alone or as part of the artifacts of the ANA Profiler Tool (available from the Cisco Developer Network).	n/a
--change_backup_time	Allows to change the database backup time.	Change the Embedded Database Backup Schedule, page 2-15
--set_smtp_server	Changes the SMTP server for e-mail notifications from the database.	This topic for an example.

You must be logged in as *pnuser* to use this command.

The following illustrates how to use the start and stop options:

```
# emdbctl --stop
Stopping Prime Network
Stopping NCCM DM Server...
- DM server is up, about to shut it down
- Sent graceful shutdown command to the dm Server (pid 25499), waiting for 2 seconds
- Checking if DM server is still up (1st)
- The DM Server is down
AVM unregistered successfully
Stopping AVMs.....Done.
Stopping the database and listener
#
# emdbctl --start

- Starting the database and listener
- Starting MVM..........................................................Done.
- Starting Gateway ..............................................................Done.
```

Add Storage to an Embedded Database

Prime Network provides two utilities for adding additional storage to an embedded database:

- To add storage to the entire database, see Add Datafiles to the Database (add_emdb_storage.pl), page 8-20.
- To add storage to a specific tablespace, see Add Datafiles to a Specific Tablespace (add_storage_for_tablespace.pl), page 8-21.

Add Datafiles to the Database (add_emdb_storage.pl)

Use the **add_emdb_storage.pl** script to add database files according to the database size you estimate you will need. When you use these scripts you will be prompted to enter your database profile (the estimated database capacity) and the history size for events and workflows. This enables the script to calculate the maximum size of the database, and to create the data files, temp files, and redo logs.

If you need assistance estimating the database size, contact your Cisco representative. The representative can provide the Memory Assessment Tool to help you with the sizing.

Step 1 Log into the Prime Network gateway as *pnuser*.

Step 2 Change directories to *NETWORKHOME*/Main/scripts/embedded_db and enter the following command:

```
# ./add_emdb_storage.pl
```

Step 3 Enter the appropriate response at the prompts:

```
- writing log to /export/home/pn310/Main/logs/emdb/add-storage-xxx.log
- Retrieving registry information & initializing connection
- How would you estimate your DB profile?
-------------------------------------------------------------------------------------
1) 1 actionable events per second (POC/LAB deployment)
2) Up to 5 actionable events per second
3) Up to 20 actionable events per second
4) Up to 50 actionable events per second
5) Up to 100 actionable events per second
6) Up to 200 actionable events per second
7) Up to 250 actionable events per second
(1  - 7) [default 1]
- Insert the event archiving size in days. Prime Network default archive is 14 days:
[default 14]
- Insert the workflow archiving size in days. Prime Network default archive is 7 days:
[default 7]
```

Note If you enter incorrect values—such as the wrong database profile estimate—you can rerun the script with different inputs.

If you encounter any errors, messages similar to the following examples are displayed.

- If there is not enough disk space to create the additional database files or redo logs:

```
- There isn't enough space on the current disks to create an additional of
        6144 MB.   Please enter a new location for creating the remaining
        DB files. Before you continue:
        1. Verify user <os-db-user> has writing permissions on the new location
        or run the following command as the OS root user:
        chown -R <os-db-user>:oinstall <path>
        2. Verify the new location is mounted as UFS with 'forcedirectio'
        option

New location:
```

Enter another location.

- If the files or redo logs cannot be created for any reason, you will see an error message and the following prompt:

```
- How would you like to continue?
---------------------------------
1) Retry
2) Skip (move to the next in list)
3) Abort
 (1 - 3) [default 1]
```

For example, if the correct permissions were not set, you would see the following.

```
Failed to add datafile for pn310:
-1119: ORA-01119: error in creating database file '/2del/pn310_DATA11.dbf'
ORA-27040: file create error, unable to create file
Linux-x86_64 Error: 13: Permission denied
```

The menu choices provide you with an opportunity to fix the permissions and retry creating the file or log.

The log file is located in *NETWORKHOME*/Main/logs/emdb/add-storage-*time-stamp*.log.

Add Datafiles to a Specific Tablespace (add_storage_for_tablespace.pl)

Use the **add_storage_for_tablespace.pl** script to add database files to a specific tablespace. If a tablespace exceeds its capacity, Prime Network will add a new data file to the tablespace and generate an hourly system event until the problem is fixed.

The command is located in *NETWORKHOME*/Main/scripts/embedded_db. It takes the following arguments:

add_storage_for_tablespace.pl --tablespace ***tablespace_name*** **--space** ***additional_space_required (MB)*** **--location** ***location_for_new_files***

The log file is located in *NETWORKHOME*/Main/logs/emdb/add-storage-to_tbs-*timestamp*.log.

Before You Begin

You will need the following information to use this script:

- The name of the tablespace that requires more datafiles.
- Additional space required for the above tablespace.
- The full directory name where the new datafiles will be created.

The following examples add 100 MB to the pn310 tablespace located in /export/home/oracle/oradata/anadb. This command performs the operation in one command line:

```
# ./add_storage_for_tablespace.pl --tablespace pn310 --space 100 --location
/export/home/oracle/oradata/anadb/
```

This procedure adds the tablespace using interactive mode:

Step 1 Log into the Prime Network gateway as *pnuser*.

Step 2 Change directories to *NETWORKHOME*/Main/scripts/embedded_db and enter the following command:

```
# ./add_storage_for_tablespace.pl
```

Step 3 Enter the appropriate response at the prompts:

```
This script will add an additional datafile for a certain tablespace in the DB
---------------------------------------------------------------------------------

+Retrieving registry information & initializing connection
+Choose one of the following Prime Network tablespaces to add datafiles to:

TABLESPACE_NAME                  FREE_SPACE_MB
------------------------------   -------------
UNDOTBS1                         1992.25
pn310_DWE                        1018.375
pn310_XMP                        1009.625
pn310_EP                          928.6875
pn310                             271.8125
pn310_ADMIN                        98.375
SYSAX                              37.375
SYSTEM                               6.75
USERS                                3.6875

- Enter tablespace name: pn310

+Choose one of the following locations for the new datafile/s to be created at:
/export/home/oracle/oradata/anadb/
- Enter location: /export/home/oracle/oradata/anadb/

- Enter the required size in MB (For Example: 1000): 100

+About to add 100 MB to pn310 on /export/home/oracle/oradata/anadb/
Successfully added 100 M on /export/home/oracle/oradata/anadb/ to pn310
```

Change the SMTP Server for Embedded Database Notifications

If necessary, you can change the SMTP server for e-mail notifications from the embedded database using the **embdctl** command, as shown in this example.

```
# emdbctl --set_smtp_server
Enter your SMTP server IP/Hostname: 1.1.1.1
Verifying connectivity to 1.1.1.1
              Failed to connect to 1.1.1.1 on port 25. Please try again
Enter your SMTP server IP/Hostname: outbound.cisco.com
Verifying connectivity to outbound.cisco.com
Reading Prime Network registry
Updating the SMTP server parameter in the database
Done
```

Prevent System Overloads (Advanced Overload Prevention/Safe Mode)

Prime Network uses a software mechanism called Automatic Overload Prevention (AOP) to detect and prevent system overload. The AOP service monitors the load produced by components in Prime Network. Similar components, such as those that control fault management, are grouped together into an AOP subsystem. When a subsystem's processing load becomes heavy, the whole system moves into *safe mode*. Other subsystems respond by adjusting their processing in order to prevent system overload. When this happens, a System event is generated and can be viewed in Prime Network EventVision.

If the subsystem continues to be overloaded, the components will take other measures to lessen the system load (if those measures are configured). As soon as the problematic subsystem returns to a normal load, all other components revert to normal.

The AOP mechanism is currently used by the following subsystems, due to the very large amount of data they process:

- Reporting subsystem.
- Fault subsystem, which includes the Alarm Plugin, Global Event Filter Agent, Event Integrity Agent, and Ticket Agent.

Loads and Running Levels

The AOP service maintains the following information about each component in a subsystem.

Load Indicator	Definition
Current Load	Current processing load. When a component's Current Load changes, other components may respond by changing their Current Loads and/or Running Levels. Supported Current Loads are: • **NORMAL** • **LOAD*x* (safe mode), where *x* is 1-6**
Running Level	The state in which a component is running. Running Levels can change in response to Current Load and/or Running Level changes in other components. Supported Running Levels are: • **NORMAL**, also called Running Level 0. • **AOP***x* or safe mode, where *x* is Running Levels 1-6.

When a problem occurs and a component's load increases, the following can occur, depending on your system configuration:

- The reporting subsystem disabled reports (at AOP 6, by default).
- The Alarm Plugin stops auto-clearing events (at AOP 6, by default).
- The Global Event Filter drops some syslogs and traps (it does this at all AOP levels, and at AOP 6, it drops *all* syslogs and traps).

To specify which events the fault subsystem drops at different running levels, see Configuring the AOP Global Event Filter, page 8-24.

Note Dropping syslogs and traps in this context means that syslogs and traps are not correlated and forwarded to the Fault Agent (AVM 25); syslogs and traps are still sent to the Event Archive (if enabled). Also note that *only* syslogs and traps are dropped; Service events and non-network events (Audit, Security, System, and Provisioning events) are *never* dropped by the AOP mechanism.

As soon as the load returns to normal on the problematic component, all components respond by returning to normal and the system moves out of safe mode.

Displaying Current AOP Loads and Running Levels

To display the status of all components that are using AOP:

Step 1 Open an SSH session to the Prime Network gateway server and log in as *pnuser*.

Step 2 Enter the following:

```
# telnet 0 2011
Connected to 0.
Sheer BOS AVM management
AVM11#/>cd aop
AVM11#aop>getAOPStatus
------------------------------------------------------------------------------------------
Subsystem  ComponentId                Load     Running Level  Last Modification Time
------------------------------------------------------------------------------------------
FAULT      ALARM_PLUGIN               NORMAL   AOP6           Thu Oct 21 13:20:20 PST 2011
FAULT      EVENT_GLOBAL_FILTER_AGENT  NORMAL   AOP1           Thu Oct 21 13:20:20 PST 2011
FAULT      EVENTINTEGRITY_AGENT       NORMAL   AOP1           Thu Oct 21 13:20:20 PST 2011
FAULT      TICKET_AGENT               LOAD1    NORMAL         Thu Oct 21 13:20:20 PST 2011
REPORTS    REPORTS_AGENT              NORMAL   AOP6           Thu Oct 21 13:20:17 PST 2011
------------------------------------------------------------------------------------------
total rows in report: 5
```

Configuring the AOP Global Event Filter

The Event Global Filter has two flavors:

- Filtering when the system is running in NORMAL mode (Running Level 0)
- Filtering when the system is in AOP mode (Running Levels AOP 1-6)

You can define filters for Running Levels 0-5—that is, for NORMAL mode, and for AOP 1-5. At Running Level 6, all traps and syslogs are dropped so no further filtering is useful.

The filter contains a list of rules that define what events should be excluded. Events are assigned a number (1-6), corresponding to the AOP running levels. When the AOP running level is *x*, all events with a number equal to or lesser than *x* are dropped. Note that this is done after events are sent to the Event Archive (if enabled).

Use the following procedure to create a new filter. In this procedure you will specify:

- Running level at which to drop the events that match the filter.
- The event information. When matched, the event will be dropped.

To create a new filter, use this procedure. For information on the properties described in the procedure, see the *Cisco Prime Network Integration Developer Guide*.

Step 1 Log into the Prime Network gateway server as *pnuser*.

Step 2 Add the new filter information to the registry using the following command.

ID is the AOP running level at which to drop events if they match the filter criteria, and *propertyName* is the event attribute to be checked by the filter:

```
# ./runRegTool.sh -gs 127.0.0.1 add 0.0.0.0
site/event-global-filter/runningLevelID/propertyName
```

propertyName can be any of the following:

Attribute	Description and Supported Values
SeverityEnum	An integer that represents the severity. Supported SeverityEnums are: 1—INFO 2—CLEARED 3—WARNING 4—MINOR 5—MAJOR 6—CRITICAL
Name	An integer that represents the alarm as defined in the alarm-types.xml registry file. For example, 1 represents “Link Down.”
State	Short description of the event, such as “Port down due to card down.”
DetectionType	An integer that represents the event protocol type. Supported DetectionTypes are: 0—Service Event 1—Syslog Event 2—V1 Trap 3—V2 Trap 4—V3 Trap

This command adds a SeverityEnum property value to AOP 1:

```
# ./runRegTool.sh -gs 127.0.0.1 add 0.0.0.0
site/event-global-filter/runningLevel1/SeverityEnum
```

Step 3 Set a value for the event property. Events will be dropped when the property has that value.

```
# ./runRegTool.sh -gs 127.0.0.1 set 0.0.0.0
site/event-global-filter/runningLevelID/propertyName/propertyValue “”
```

This command sets the SeverityEnum value to 1 in the Global Event Filter:

```
# ./runRegTool.sh -gs 127.0.0.1 set 0.0.0.0
site/event-global-filter/runningLevel1/SeverityEnum/1 “”
```

To remove a filter, use this procedure.

Step 1 Log into the Prime Network gateway server as *pnuser*.

Step 2 Remove the filter from the registry using the following command.

ID is the AOP running level at which to drop events if they match the filter criteria, and *propertyName* is the event attribute to be checked by the filter:

```
# ./runRegTool.sh -gs 127.0.0.1 remove 0.0.0.0
site/event-global-filter/runningLevelID/propertyName ""
```

This command removes the filter created in the previous procedure:

```
# ./runRegTool.sh -gs 127.0.0.1 remove 0.0.0.0
site/event-global-filter/runningLevel1/SeverityEnum ""
```

Track Database and System Integrity Events

The following predefined reports can provide you with important database statistics for a period of time that you specify. To run any of these reports, select **Reports > Run Report > Events Reports** and choose the report name.

Report Name	Provides the following info for a specified period:
Fault DB vs. Event Archive Statistics	Lists the total number of these types of events for a specified time period: Syslogs, traps, tickets, correlated and uncorrelated events, network events, non-network events, and service events
Database Monitoring	(includes events per second)—How many of the following occurred in the specified time period: Active tickets, active alarms, and active events; large tickets (and their event counts); notifications. You can also get the event rate per second for that period.
Detailed System Events	Lists all System events for the gateway server, such as ticket archiving, tablespace, partitioning, dropping events, and AOP events